

Functional correspondence by matrix completion

Artiom Kovnatsky¹ Michael M. Bronstein¹ Xavier Bresson²
 Pierre Vandergheynst²

¹Università della Svizzera italiana (USI), Switzerland

²Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

December 30, 2014

Abstract

In this paper, we consider the problem of finding dense intrinsic correspondence between manifolds using the recently introduced functional framework. We pose the functional correspondence problem as matrix completion with manifold geometric structure and inducing functional localization with the L_1 norm. We discuss efficient numerical procedures for the solution of our problem. Our method compares favorably to the accuracy of state-of-the-art correspondence algorithms on non-rigid shape matching benchmarks, and is especially advantageous in settings when only scarce data is available.

1 Introduction

Correspondence problems are, without doubt, among the most important problems that have to be dealt when analyzing highly-complex geometric data. Establishing correspondence between images is a cornerstone of many algorithms in computer vision and image analysis [50]. In geometry processing and graphics, finding intrinsic correspondence between deformable 3D shapes is one of the fundamental problems with a plethora of applications ranging from texture mapping to animation [16]. In machine learning applications, correspondence problems in high-dimensional spaces arise in multi-modal data analysis problems such as manifold alignment and multi-view clustering [49, 14].

Related work A wide class of correspondence methods aims at minimizing some structure distortion, which can include similarity of local features [32, 9, 2, 51, 24], geodesic [27, 7] or diffusion distances [12, 8], or a combination thereof [45]. Zeng et al. [53] used higher-order structures. Typically, the computational complexity of such methods is high, and there have been several attempts to alleviate the computational complexity using hierarchical [38] or smart sampling [44] methods. Several approaches formulate the correspondence problem as an NP-hard quadratic assignment and propose different relaxations thereof [46, 21, 36]. A recent trend is to use machine learning techniques, such as random forests, to learn correspondences [39, 37].

Another important class of methods strives to embed the intrinsic structure of the manifold into some other space, where the correspondence can be parametrized with a small number of degrees of freedom. Elad and Kimmel [13] used multi-dimensional scaling to embed the geodesic metric of the matched manifolds into the Euclidean space, where alignment of the resulting “canonical forms” is then performed by simple rigid matching (ICP) [10, 3, 29]. Mateus et al. [25] used the first eigenfunctions of the Laplace-Beltrami operator as embedding coordinates. A more recent method of spectral kernel map [40] follows this path. Lipman et al. [23, 18, 19] used conformal embeddings to parametrize correspondence as a Möbius transformation.

More recently, there is an emerging interest in *soft correspondence* approaches. Several methods formulated soft correspondence as a mass-transportation problem [26, 42]. Ovsjanikov et al. [31] introduced the *functional correspondence* framework, modeling the correspondence as a linear operator between spaces of functions on two manifolds, which has an efficient representation in the Laplacian eigenbases. In such a formulation, the problem of finding correspondence is reduced to solving a linear system of equations given some known set of corresponding functions on the two manifolds. Pokrass et al. [35] extended this approach to the case when the ordering of the corresponding functions is unknown, and introduced a regularization on the diagonal structure of the correspondence matrix representation in the Laplacian eigenbases. Kovnatsky et al. [20] proposed constructing coupled bases by simultaneous diagonalization of Laplacians. In such bases the correspondence is represented by an approximately diagonal matrix, which allows to solve the linear system of equations for diagonal elements only.

Main contribution Our paper deals with the problem of finding dense intrinsic correspondence between manifolds in the functional formulation. We propose treating functional correspondence as *geometric matrix completion*. Our approach is inspired by the recent work on recovery of matrices on graphs [17], which introduced geometric structure into the matrix completion problem. Treating the rows and columns of the functional correspondence matrix as vector-valued functions on the respective manifolds, we introduce geometric structure using the Dirichlet energy. We show that our method includes the previous approaches of [31, 20] as particular settings, and compares favorably to state-of-the-art methods for non-rigid shape correspondence on the challenging Princeton benchmark [19]. The advantage of our method is especially pronounced in the *scarce data* settings. The proposed method is completely generic and can be applied to any manifolds and high-dimensional geometric data, such as visual data and annotated images.

2 Background

Notation In this paper, we use bold capital letters to denote matrices, bold lower-case letters for vectors, and italic lower-case letters for scalars. Given a matrix \mathbf{A} , we denote by $\|\mathbf{A}\|_F = \left(\sum_{ij} a_{ij}^2\right)^{1/2}$ its Frobenius norm, by $\|\mathbf{A}\|_1 = \sum_{ij} |a_{ij}|$ its L_1 -norm, and by $\|\mathbf{A}\|_* = \sum_i \sigma_i$ its nuclear (or trace) norm, where $\sigma_1, \sigma_2, \dots$ denote the singular




Figure 1: Laplacian construction for a graph (left) and a triangular mesh (right).

values of \mathbf{A} . Alternatively, provided a decomposition of the form $\mathbf{A} = \mathbf{U}\mathbf{V}^\top$, the nuclear norm can be written as $\|\mathbf{A}\|_* = \frac{1}{2}(\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2)$ [43]. We denote by δ_i the Kronecker delta, i.e. a unit vector with one at position i and all the rest zeros.

Laplacians Let us be given a Riemannian manifold sampled at n points $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subseteq \mathbb{R}^d$. The local structure of the manifold is modeled as an undirected graph (X, E) in case of abstract data, or as a simplicial complex (triangular mesh) (X, E, F) in case of 3D shapes, where $E \subseteq X \times X$ and $F \subset X^3$ denote the edges and faces, respectively. We denote by $L^2(X) \cong \mathbb{R}^n$ the space of real functions $\mathbf{f} = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_n))$ on the discrete manifold.

We define the *Laplacian* of X as $\mathbf{L} = \mathbf{A}^{-1}(\mathbf{D} - \mathbf{W})$, where $\mathbf{D} = \text{diag}(\sum_{k \neq i} w_{ki})$, and \mathbf{W} and $\mathbf{A} = \text{diag}(a_1, \dots, a_n)$ are some edge- and vertex weights, respectively. The definition of the weights depends on the particular discretization of the manifold. In this paper, we consider two types of objects (see Figure 1):

Graphs, for which we use the Gaussian edge weights

$$w_{ij} = \begin{cases} e^{-\|\mathbf{x}_i - \mathbf{x}_j\|^2 / 2\sigma^2} & ij \in E; \\ 0 & \text{else,} \end{cases} \quad (1)$$

and construct the edge set E using K -nearest neighbors. The vertex weight is defined as $\mathbf{A} = \mathbf{D}$ (random walk Laplacian) or $\mathbf{A} = \mathbf{I}$ (unnormalized Laplacian) [48]. 3D shapes represented as point clouds are treated in this way.

Triangular meshes, for which we use the discretization of the Laplace-Beltrami operator [22] based on the classical cotangent weights [34, 28]

$$w_{ij} = \begin{cases} (\cot \alpha_{ij} + \cot \beta_{ij})/2 & ij \in E; \\ 0 & \text{else,} \end{cases} \quad (2)$$

where α_{ij}, β_{ij} are the angles in from of the edge ij . The vertex weights are defined as local area elements $a_i = \frac{1}{3} \sum_{j,k: ij \in F} \text{area}(\mathbf{x}_i \mathbf{x}_j \mathbf{x}_k)$, equal to one third of the sum of the one-ring triangles areas (marked in green in Figure 1).

Harmonic analysis on manifolds The eigenvectors $\Phi = (\phi_1, \dots, \phi_n)$ of the Laplacian satisfying $\mathbf{L}\Phi = \Phi\Lambda$ (here $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ are the corresponding eigenvalues and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$) form an orthonormal basis that generalizes the Fourier basis to non-Euclidean spaces. Given a function $\mathbf{f} \in L^2(X)$, its Fourier coefficients are computed as $\alpha = \Phi^\top \mathbf{f}$. Taking the first k coefficients, one obtains a smooth (“low-pass”) approximation of the function $\mathbf{f} \approx \sum_{l=1}^k \alpha_l \phi_l = \Phi_k \Phi_k^\top \mathbf{f}$, where $\Phi_k = (\phi_1, \dots, \phi_k)$.

Compressed manifold modes The Laplacian eigenbasis can be found by minimizing the *Dirichlet energy*,

$$\min_{\Phi} \quad \text{tr}(\Phi^\top \mathbf{L} \Phi) \quad \text{s.t.} \quad \Phi^\top \Phi = \mathbf{I}. \quad (3)$$

Neumann et al. [30], following Ozoliņš et al. [33], proposed constructing localized approximate eigenbases on manifolds (referred to as *compressed manifold modes* or CMM) by imposing the L_1 -penalty on $\|\Phi\|_1$, which, in combination with the Dirichlet energy makes the basis functions localized,

$$\min_{\Phi} \quad \text{tr}(\Phi^\top \mathbf{L}_X \Phi) + \mu \|\Phi\|_1 \quad \text{s.t.} \quad \Phi^\top \Phi = \mathbf{I}, \quad (4)$$

for $\mu > 0$.

Functional maps Let X and Y denote two manifolds sampled at n and m points, respectively. Ovsjanikov et al. [31] propose to model the *functional correspondence* between the spaces $L^2(X)$ and $L^2(Y)$ as the $m \times n$ matrix \mathbf{T}_0 , which maps a function $\mathbf{f} \in L^2(X)$ into $\mathbf{T}_0 \mathbf{f} = \mathbf{g} \in L^2(Y)$. Traditional point-wise correspondence is a particular case of this model wherein \mathbf{T}_0 maps delta functions into delta functions.

Let \mathbf{L}_X and \mathbf{L}_Y be the Laplacians of X and Y , and let $\Phi_k = (\phi_1, \dots, \phi_k)$ and $\Psi_k = (\psi_1, \dots, \psi_k)$ be the respective truncated Laplacian eigenbases. Let us be given q corresponding functions $\mathbf{F} = (\mathbf{f}_1, \dots, \mathbf{f}_q)$ and $\mathbf{G} = (\mathbf{g}_1, \dots, \mathbf{g}_q)$ satisfying $\mathbf{T}_0 \mathbf{F} = \mathbf{G}$, where \mathbf{T}_0 is the unknown groundtruth correspondence. The functional correspondence can be approximated in these bases as a rank- k matrix $\mathbf{T} = \Psi_k \mathbf{C} \Phi_k^\top$, where \mathbf{C} is a $k \times k$ matrix translating Fourier coefficients from the basis Φ_k to the basis Ψ_k .

The matrix \mathbf{C} can be found by solving the system of equations $\mathbf{G} = \Psi_k \mathbf{C} \Phi_k^\top \mathbf{F}$ in the least-squares sense,

$$\min_{\mathbf{C}} \|\mathbf{C} \Phi_k^\top \mathbf{F} - \Psi_k^\top \mathbf{G}\|_{\mathbf{F}}^2. \quad (5)$$

Note that since the groundtruth correspondence \mathbf{T}_0 is unknown, in practice the matrices \mathbf{F}, \mathbf{G} must be computed *independently* on X and Y such that $\mathbf{T}_0 \mathbf{F} \approx \mathbf{G}$. In the simplest case, $\mathbf{F} = (\delta_{i_1}, \dots, \delta_{i_q})$, $\mathbf{G} = (\delta_{j_1}, \dots, \delta_{j_q})$ represent some known point-wise correspondences (“seeds”) between points $\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_q}$ and $\mathbf{y}_{j_1}, \dots, \mathbf{y}_{j_q}$. In shape correspondence applications, \mathbf{F}, \mathbf{G} are typically computed using some intrinsic shape descriptor such as HKS [32], WKS [2], MeshHOG [51], ShapeMSER [24], which is invariant to shape deformations.

For applications requiring point-wise maps, Ovsjanikov et al. [31] devise an iterative conversion scheme similar to iterative closest point (ICP) methods [10, 3, 29]. Thinking of Φ_k and Ψ_k as k -dimensional point clouds (with n and m points, respectively), the matrix C can be interpreted of as a rigid alignment between them. Starting with some initial matrix C , first find the closest row j_i in $\Psi_k C$ for each i th row of Φ_k (this step is equivalent to closest point correspondence in ICP). Second, find an orthonormal C minimizing $\sum_i \|(\Phi_k)_i - (\Psi_k C)_{j_i}\|$ (this is equivalent to the alignment step in ICP). The process is repeated until convergence, producing point-wise correspondences $(\mathbf{x}_1, \mathbf{y}_{j_1}), \dots, (\mathbf{x}_n, \mathbf{y}_{j_n})$.

Finally, note that in order for system (5) to be (over-) determined, we must have $k \leq q$, where q is the number of given corresponding functions (data). At the same time, the quality of the correspondence depends on k : the more basis elements are taken, the better, since due to truncation of the Fourier expansion, small k results in poor spatial localization of the correspondence. Furthermore, truncated Fourier series of a discontinuous function manifest oscillations, a behavior known in harmonic analysis as the ‘‘Gibbs phenomenon’’ (see examples in Figure 5).

Permuted sparse coding Pokrass et al. [35] addressed the case when the ordering of the columns of F, G is unknown and is modeled by a $k \times k$ permutation matrix Π . Furthermore, they noticed that for near-isometric manifolds, the eigenbases satisfy $\phi_i^\top T_0^\top \psi_j \approx \pm \delta_{ij}$, i.e., the matrix C is approximately diagonal. The diagonal structure of C was induced using an L_1 penalty,

$$\min_{\Pi, C} \|C\Phi_k^\top F\Pi - \Psi_k^\top G\|_F^2 + \mu\|M \circ C\|_1, \quad (6)$$

where $\mu > 0$ and M is a weight matrix with zero diagonal and \circ denotes the Hadamard (element-wise) matrix multiplication. The problem was solved by means of alternating minimization w.r.t. Π (which turns out a linear assignment problem) and C (which is equivalent to sparse coding).

Coupled diagonalization Kovnatsky et al. [20] noticed that the matrix C encoding the correspondence depends on the choice of bases, and proposed finding new approximate eigenbases $\hat{\Phi}_k = \Phi_k P$ and $\hat{\Psi}_k = \Psi_k Q$ by means of $k \times k$ orthonormal matrices P, Q in which the Fourier coefficients of F and G are as similar as possible. The new bases behave as eigenbases if they approximately diagonalize the respective Laplacians, i.e., $\hat{\Phi}_k^\top L_X \hat{\Phi}_k = P^\top \Phi_k^\top L_X \Phi_k P = P^\top \Lambda_{X,k} P$ (respectively, $\hat{\Psi}_k^\top L_Y \hat{\Psi}_k = Q^\top \Lambda_{Y,k} Q$), where $\Lambda_{X,k} = \text{diag}(\lambda_1^X, \dots, \lambda_k^X)$ and $\Lambda_{Y,k} = \text{diag}(\lambda_1^Y, \dots, \lambda_k^Y)$. This property is enforced in the optimization problem

$$\begin{aligned} \min_{P, Q} \quad & \|F^\top \Phi_k P - G^\top \Psi_k Q\|_F^2 + \mu_1 \text{off}(P^\top \Lambda_{X,k} P) \\ & + \mu_2 \text{off}(Q \Lambda_{Y,k} Q^\top) \quad \text{s.t.} \quad P^\top P = Q^\top Q = I, \end{aligned} \quad (7)$$

(with $\mu_1, \mu_2 > 0$) using the off-diagonal penalty $\text{off}(A) = \sum_{i \neq j} a_{ij}^2$.

After P, Q are found, in the new bases $\Phi_k P, \Psi_k Q$ the matrix C is approximately diagonal. Then, the system of equations $C P^\top \Phi_k^\top F = Q^\top \Psi_k^\top G$ is solved for diagonal elements of C only.




Figure 2: Illustration of functional correspondence as matrix completion. Column $\mathbf{t}_i = \mathbf{T}\delta_i$ of matrix \mathbf{T} is the functional map of a delta at \mathbf{x}_i . Geometric structure is imposed on the columns to ensure smoothness, i.e. $\mathbf{t}_i, \mathbf{t}_{i'}$ corresponding to spatially close points $\mathbf{x}_i, \mathbf{x}_{i'}$ on X are similar.

3 Functional maps as matrix completion

Kalofolias et al. [17] studied the problem of *matrix completion on graphs*, where the rows and columns of the matrix have underlying geometric structure. They show that adding geometric structure to the standard matrix completion problem improves recovery results.

We use the same philosophy to formulate the problem of finding a functional map as matrix completion, whose rows and columns are interpreted as functions on the respective manifolds X and Y . For this purpose, we consider the matrix \mathbf{T} as a collection of columns $\mathbf{T} = (\mathbf{t}_1, \dots, \mathbf{t}_n)$ or rows $\mathbf{T} = (\mathbf{t}^1, \dots, \mathbf{t}^m)^\top$, where \mathbf{t}_i and \mathbf{t}^j denote the i th column and j th row of \mathbf{T} , respectively. The column $\mathbf{t}_i = \mathbf{T}\delta_i$ is the function on Y corresponding to a delta located at point \mathbf{x}_i on X . Similarly, the row $\mathbf{t}^j = \delta_j^\top \mathbf{T}$ is the function on X corresponding to a delta located at point \mathbf{y}_j on Y .

As in the classical matrix completion problem, we aim at recovering the unknown correspondence matrix \mathbf{T}_0 from a few observations of the form $\mathbf{T}_0 \mathbf{F} = \mathbf{G}$, looking for a matrix $\mathbf{T} \approx \mathbf{T}_0$ that explains the data in a “simple” way, in the sense discussed in the following. Our problem comprises the following terms:

Data term The correspondence should respect the data, which is achieved by minimizing $\|\mathbf{T}\mathbf{F} - \mathbf{G}\|_F$.

Smoothness The correspondence must be *smooth*, in the sense that if $\mathbf{x}_i, \mathbf{x}_{i'}$ are two close points on X , then the respective corresponding functions are similar, i.e., $\mathbf{t}_i \approx \mathbf{t}_{i'}$ (see Figure 2). Similarly, for close points $\mathbf{y}_j, \mathbf{y}_{j'}$ on Y , the rows $\mathbf{t}^j \approx \mathbf{t}^{j'}$. Smoothness is achieved by minimizing the *row* and *column Dirichlet energy*

$$\text{tr}(\mathbf{T}^\top \mathbf{L}_Y \mathbf{T}) + \text{tr}(\mathbf{T} \mathbf{L}_X \mathbf{T}^\top).$$

Localization The correspondence is localized using the L_1 -penalty $\|\mathbf{T}\|_1$, which, in combination with smoothness, results in a few non-zero elements that are close in space. Note that differently from CMM [33, 30] where the penalty is applied to the basis functions, we apply it to the correspondence matrix.

“Simplicity” By simplicity, we mean here that the correspondence matrix is ‘explained’ using a small number of degrees of freedom. The following models are commonly used in matrix completion and recommendation systems literature.

Fixed rank. The most straightforward model is to assume that $\text{rank}(\mathbf{T}) \leq k \ll m, n$. In this setting, we can decompose $\mathbf{T} = \mathbf{U}\mathbf{V}^\top$ into the factors \mathbf{U}, \mathbf{V} of size $m \times k$ and $n \times k$, respectively. Combining all the above terms leads to the optimization problem

$$\begin{aligned} \min_{\mathbf{U}, \mathbf{V}} \quad & \|\mathbf{U}\mathbf{V}^\top \mathbf{F} - \mathbf{G}\|_F^2 + \mu_1 \text{tr}(\mathbf{U}\mathbf{V}^\top \mathbf{L}_X \mathbf{V}\mathbf{U}^\top) \\ & + \mu_2 \text{tr}(\mathbf{V}\mathbf{U}^\top \mathbf{L}_Y \mathbf{U}\mathbf{V}^\top) + \mu_3 \|\mathbf{U}\mathbf{V}^\top\|_1, \end{aligned} \quad (8)$$

where $\mu_1, \mu_2, \mu_3 > 0$ are parameters determining the tradeoff between smoothness and localization.

Low rank. Another popular model is to find the matrix with smallest $\text{rank}(\mathbf{T})$. However, this minimization is known to be NP-hard, and the nuclear norm $\|\mathbf{T}\|_*$ is typically used as a convex proxy, leading to

$$\begin{aligned} \min_{\mathbf{T}} \quad & \|\mathbf{T}\mathbf{F} - \mathbf{G}\|_F^2 + \mu_1 \text{tr}(\mathbf{T} \mathbf{L}_X \mathbf{T}^\top) + \mu_2 \text{tr}(\mathbf{T}^\top \mathbf{L}_Y \mathbf{T}) \\ & + \mu_3 \|\mathbf{T}\|_1 + \mu_4 \|\mathbf{T}\|_*. \end{aligned} \quad (9)$$

This problem is convex and is typically solved using augmented Lagrangian methods [5].

Low norm. Srebro et al. [43] rewrite problem (9) using the decomposition $\mathbf{T} = \mathbf{U}\mathbf{V}^\top$ with \mathbf{U} and \mathbf{V} of size $m \times k$ and $n \times k$, respectively. Note that unlike the fixed rank case, here k can be arbitrarily large. The nuclear norm is written as $\|\mathbf{T}\|_* = \frac{1}{2}(\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2)$. We thus have the problem

$$\begin{aligned} \min_{\mathbf{U}, \mathbf{V}} \quad & \|\mathbf{U}\mathbf{V}^\top \mathbf{F} - \mathbf{G}\|_F^2 + \mu_1 \text{tr}(\mathbf{U}\mathbf{V}^\top \mathbf{L}_X \mathbf{V}\mathbf{U}^\top) \\ & + \mu_2 \text{tr}(\mathbf{V}\mathbf{U}^\top \mathbf{L}_Y \mathbf{U}\mathbf{V}^\top) + \mu_3 \|\mathbf{U}\mathbf{V}^\top\|_1 \\ & + \frac{\mu_4}{2} (\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2). \end{aligned} \quad (10)$$

Unlike (9), this problem is non-convex w.r.t. both \mathbf{U} and \mathbf{V} , however behaves well for large k [43].

3.1 Subspace parametrization

The main disadvantage of problems (8)–(10) is that the number of variables depends on the number of samples m, n , which may result in scalability issues for large $(m, n \sim$

10^6) manifolds. To overcome this issue, we approximate the $n \times k$ and $m \times k$ factors \mathbf{U}, \mathbf{V} in the truncated Laplacian eigenbases of X and Y using $k' \geq k$ first expansion terms, $\mathbf{U} \approx \Psi_{k'} \mathbf{A}$ and $\mathbf{V} \approx \Phi_{k'} \mathbf{B}$, where matrices \mathbf{A}, \mathbf{B} of the expansion coefficients are of size $k' \times k$. This leads to a subspace version of problem (10),

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{B}} \quad & \|\mathbf{A}\mathbf{B}^\top \Phi_{k'}^\top \mathbf{F} - \Psi_{k'}^\top \mathbf{G}\|_F^2 \\ & + \mu_1 \text{tr}(\mathbf{A}\mathbf{B}^\top \Lambda_{X,k'} \mathbf{B}\mathbf{A}^\top) + \mu_2 \text{tr}(\mathbf{B}\mathbf{A}^\top \Lambda_{Y,k'} \mathbf{A}\mathbf{B}^\top) \\ & + \mu_3 \|\Psi_{k'} \mathbf{A}\mathbf{B}^\top \Phi_{k'}^\top\|_1 + \frac{\mu_4}{2} (\|\Psi_{k'} \mathbf{A}\|_F^2 + \|\Phi_{k'} \mathbf{B}\|_F^2) \end{aligned} \quad (11)$$

where we used the invariance of the Frobenius norm to orthogonal transformations and the fact that $\Phi_{k'}^\top \mathbf{L}_X \Phi_{k'} = \Lambda_{X,k'}$, $\Psi_{k'}^\top \mathbf{L}_Y \Psi_{k'} = \Lambda_{Y,k'}$ to simplify the expressions.

Note that now the number of variables $2kk'$ is *independent of the number of samples*. We emphasize that k', k can be arbitrarily large and are dictated only by complexity considerations and not by the amount of data. This is one of the major advantages of our approach compared to [31, 35, 20], which we demonstrate experimentally in the following section (see examples in Figure 5). Finally, we observe that the subspace version of the fixed rank problem (8) is a particular setting of (11) with $\mu_4 = 0$.

3.2 Relation to other approaches

Ovsjanikov et al. [31] solve a fixed-rank approximation problem for \mathbf{T} expressed in the truncated Laplacian eigenbases as $\mathbf{T} = \Psi_{k'} \mathbf{C} \Phi_{k'}$. It boils down to a particular setting of our formulation (11) with $\mathbf{C} = \mathbf{A}\mathbf{B}^\top$ and $\mu_1 = \mu_2 = \mu_3 = \mu_4 = 0$. Note that in (5), $\text{rank}(\mathbf{C}) = k$ is bounded by the size of the data q : if $k > q$, the system of equations in (5) becomes underdetermined. As a consequence, in scarce data settings ($q \ll k$), the methods performs very poorly (see Figure 5). Pokrass et al. [35] solve the problem of [31] with additional prior on the diagonal structure of \mathbf{C} , which holds only for approximately isometric manifolds. Our method is generic and can handle highly non-isometric manifolds.

Kovnatsky et al. [20] look for a pair of bases in which one tries to achieve a diagonal \mathbf{C} . The bases are constructed as a rotation of the Laplacian eigenbases (7) by means of orthonormal matrices \mathbf{A}, \mathbf{B} . This problem is equivalent to the setting of our problem (11) with $\mu_1 = \mu_4 = 0$, with the only difference that the Dirichlet energy terms are replaced with off-diagonality terms.

Finally, we emphasize that our method is “*basis-free*”: we work directly with the correspondence matrix, and rather than imposing some structure in the correspondence representation in some basis (as done in prior works of [35, 20, 30]), impose structure on the correspondence matrix *directly*. This also allows us to work simultaneously with very different manifold representations (such as meshes and graphs, see Figure 8). The factors \mathbf{U}, \mathbf{V} in our problem can be interpreted as *ad hoc* bases in which the correspondence is represented as the identity matrix, $\mathbf{T} = \mathbf{U}\mathbf{I}\mathbf{V}^\top$. These bases are coupled (see Figure 3, bottom), and thus our problem also acts similarly to the simultaneous approximate Laplacian diagonalization procedure [20].




Figure 3: A few vectors from the standard Laplacian eigenbases Φ, Ψ (first and second rows) and *ad hoc* bases \mathbf{U}, \mathbf{V} (third and fourth rows) constructed as a result of our optimization over the functional correspondence between the human and gorilla shapes. Note that our bases are coupled, i.e. $\mathbf{T}_0 \mathbf{U} \approx \mathbf{V}$. Here and in the following, hot colors represent positive values; cold colors represent negative values.

3.3 Numerical optimization

Manifold optimization We used manifold optimization toolbox `manopt` [4] to solve our problem, in which the non-differentiable L_1 -norm term was approximated using $|t| \approx \sqrt{t^2 + \xi}$. The main idea of manifold optimization is to model the space of fixed-rank matrices as a \mathcal{C}^∞ Riemannian manifold $\mathcal{M}_k = \{\mathbf{X} \in \mathbb{R}^{m \times n} : \text{rank}(\mathbf{X}) = k\}$ and perform descent (e.g. by the method of conjugate gradients) on this manifold. The intrinsic gradient of the cost function on \mathcal{M}_k is a vector in the tangent space; in order to construct the conjugate gradient search direction, gradients from previous iterations must be brought to the same tangent space by means of parallel transport on \mathcal{M}_k . After performing a descent step on the tangent space, the iterate is retracted to the manifold by means of the exponential map. For additional details on the use of manifold optimization in matrix completion problems, we refer the reader to [47].

Initialization Assume we are given some \mathbf{T} as initial correspondence matrix. We initialize our problem by first computing the SVD of $\mathbf{T} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ and then setting $\mathbf{A} = \mathbf{\Phi}_{k'}^\top \mathbf{U}\mathbf{\Sigma}^{1/2}$ and $\mathbf{B} = \mathbf{\Psi}_{k'}^\top \mathbf{V}\mathbf{\Sigma}^{1/2}$. In our experiments, we used the correspondence of Ovsjanikov et al. as the initial \mathbf{T} .

4 Results

4.1 3D Shapes

Data We used the TOSCA [6] dataset containing CAD models of human and animal with 5K-50K vertices, and the SCAPE [1] dataset, containing meshes of scanned humans in different poses with about 10K vertices. For meshes, we used the cotangent formula (2) to construct the Laplacian. Point clouds were created by removing the mesh structure from triangular meshes; 10 nearest neighbors were then used to construct unnormalized graph Laplacians according to (1) with σ found using the self-tuning method [52]. In our experiments, as the data term we used one of the following three types of descriptors (computed fully automatically):

Segments produced by the persistence-based segmentation of [41], represented as binary indicator functions. A typical number of segments per shape was 3 – 4.

Localized WKS+HKM. 100 Wave kernel signatures (WKS) [2] and 100 Heat kernel maps (HKM) [32], localized on the aforementioned segments, resulting in a total of 200 functions per segment, or 600 – 800 functions per shape.

Shape MSER. Stable regions detected using the method of [24], around 8 – 15 regions per shape. The MSER regions were represented as binary indicator functions.

The Segments, WKS, and HKM descriptors were provided by Ovsjanikov. ShapeMSER descriptors were computed using the code and settings of [24].




Figure 4: Examples of descriptors used as data (columns of \mathbf{F} , \mathbf{G}) in our experiments.

Performance criteria For the evaluation of the correspondence quality, we used the Princeton benchmark protocol [19] with two error criteria:

Hard error is the criterion used in [19] for point-wise maps. Assume that a correspondence algorithm produces a pair of points $(\mathbf{x}_{i_1}, \mathbf{y}_{j'_1})$, whereas the groundtruth correspondence is $(\mathbf{x}_{i_1}, \mathbf{y}_{j_1})$. Then, the inaccuracy of the correspondence is measured as

$$\epsilon_{\text{hard}}(x_{i_1}) = \frac{d_Y(\mathbf{y}_{j_1}, \mathbf{y}_{j'_1})}{\left(\sum_{j=1}^m a_j^Y\right)^{1/2}}, \quad (12)$$

and has units of normalized length on Y (ideally, zero). Here a^Y is the local area element on Y .

Soft error is the generalization of (12) to soft maps. Under functional correspondence, $\mathbf{T}\delta_{i_1} = \mathbf{t}_{i_1} \in L^2(Y)$, which ideally should be δ_{j_1} . The inaccuracy of functional correspondence is measured as the average geodesic distance from the j_1 th point on Y , weighted by the absolute value of the corresponding function,

$$\epsilon_{\text{soft}}(x_{i_1}) = \frac{\sum_{j=1}^m d_Y(\mathbf{y}_{j_1}, \mathbf{y}_j) |t_{j,i_1}|}{\left(\sum_{j=1}^m a_j^Y\right)^{1/2} \sum_{j=1}^m |t_{j,i_1}|}. \quad (13)$$

Note that (12) is a particular setting of (13) where $t_{j,i_1} = \delta_{j-j'_1}$.

For a set of values $\rho \geq 0$, we plot the percentage of points $\frac{1}{n}|\{x \in X : \epsilon(x) \leq \rho\}|$ producing correspondence better than ρ (the higher the curve, the better).

Influence of parameters To study the behavior of our problem, we conducted a set of experiment in ‘sterile’ conditions, using a pair of near-isometric human shapes with ideally corresponding \mathbf{F}, \mathbf{G} . The number of corresponding functions in the data term was $q = 50$. The correspondence quality was evaluated using the soft criterion (13), to rule out any artifacts introduced by conversion to point-wise map.

Rank. We set problem (11) in the fixed rank regime with parameters $k' = 350$, $\mu_1 = \mu_2 = 10^{-8}$, $\mu_3 = 10^{-1}$, $\mu_4 = 0$, $\xi = 10^{-3}$ and varied the rank k . For comparison, we ran the methods of Ovsjanikov et al. [31] and Pokrass et al. [35] with the same k . Note that the performance of [31, 35] degrades when $k > q$, as opposed to increase in the performance of our approach (Figure 5). Practically, this means that the correspondence quality they can obtain is bounded by the amount of given data, while our method can work well with very scarce data.

Localization We repeated the same experiments with $\mu_1 = \mu_2 = 10^{-4}$, $k = 45$ fixed, and varying μ_3 . Increasing the weight of the L_1 term results in better localization of the correspondence (Figure 6), and, as a consequence, better performance.

Princeton benchmark We used the Princeton evaluation protocol (allowing for symmetric matches, since our method is fully intrinsic and thus produces correspondence defined up to intrinsic symmetry) on the SCAPE dataset and followed verbatim the settings of [31, 35].

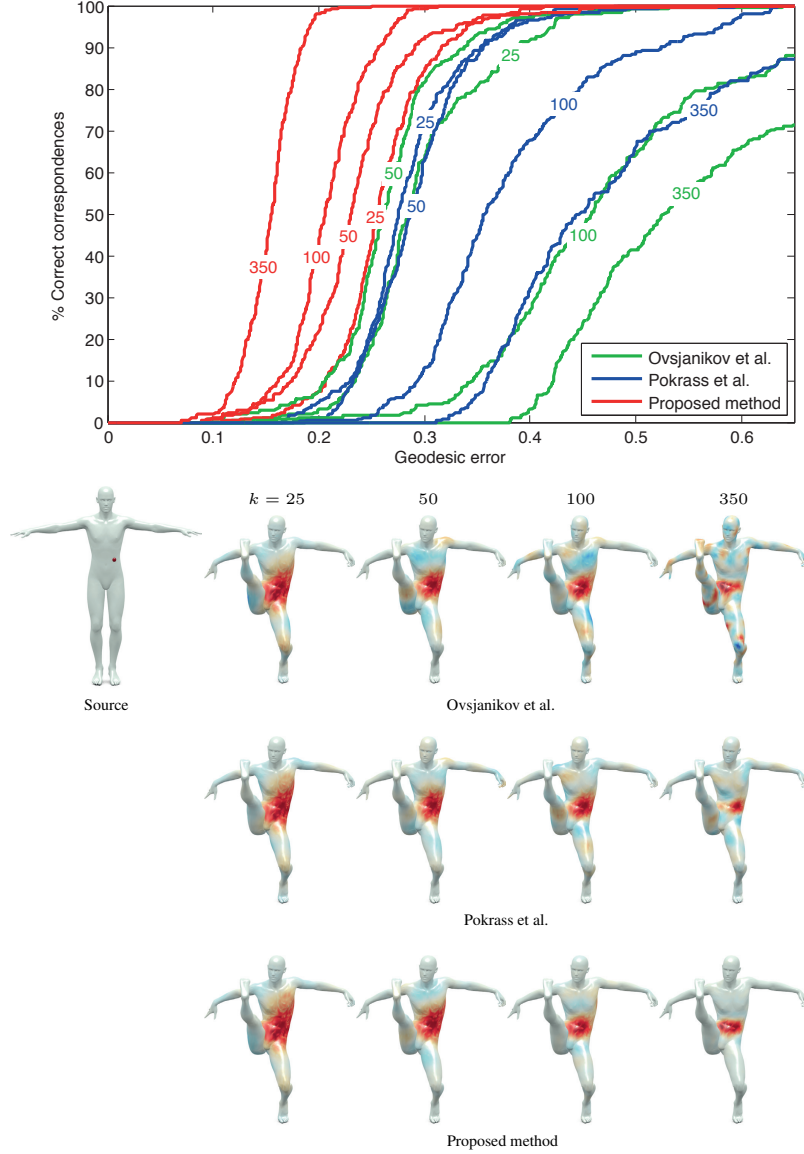


Figure 5: Influence of the rank. Top: quality (soft error) of correspondence between near isometric poses of the human shape obtained using different rank k . Note that increasing the rank above the number of corresponding functions q results in correspondence quality degradation in the methods of Ovsjanikov et al. and Pokrass et al., while with our method the performance increases. Bottom: functional correspondence of a delta function (red dot on the leftmost shape) using different methods. Note the Gibbs phenomenon manifested as negative values.




Figure 6: Influence of L_1 on the localization of the correspondence. Top: quality (soft error) of correspondence between near isometric poses of the human shape obtained using different values of μ_3 . Bottom: functional correspondence of a delta function (red dot on the leftmost shape) for different μ_3 .

We compared the methods of [31, 35] (using code provided by the authors) and our method, in two settings: *scarce data*, using Segments (one binary indicator function per segment), and *rich data*, using localized WKS+HKM (200 functions per segment). Our method was used with the parameters $k' = 350$, $k = 250$ for the first case, and $k' = 50$, $k = 30$ for the second. In both cases, we used $\mu_1 = \mu_2 = 10^{-8}$, $\mu_3 = 10^{-5}$, $\mu_4 = 10^{-8}$, and $\xi = 10^{-3}$. Resulting functional maps were converted into point-wise maps using the ICP-like method of [31]. Correspondence quality was evaluated using the hard error criterion. For comparison, we also reproduce the performance of some recent state-of-the-art point-wise correspondence algorithms: blended maps [19], best conformal (least-distorting map without blending), Möbius voting [23], HKM [32], and GMDS [7]. The results are shown in Figure 7. Our method achieves state-of-the-art performance; its advantage is especially prominent in the scarce data setting.




Figure 7: Comparison of different algorithms using the Princeton benchmark (symmetric correspondences, hard error criterion) on the SCAPE dataset. Solid line corresponds to the rich data setting (200 functions per segment); dashed line corresponds to the scarce data setting (one function per segment).

Robustness Figure 8 depicts examples of shape correspondence obtained with the proposed technique. The top row shows examples of SCAPE shape matching using Segments data. Given that this data is very scarce (3 – 4 segments per shape), we find it quite remarkable that our method is able to obtain very high quality correspondence. Second row shows examples of correspondence in the presence of geometric and topological noise, using ShapeMSER data. Our method handles well even large missing parts. Third row shows examples of non-isometric shape matching, and fourth row shows a particularly challenging setting of mesh to point cloud matching.

4.2 Abstract manifolds

Visual manifolds We used the dataset of [15] containing 831 120×100 images of a human face depicted with varying pitch and yaw angles (Figure 9, left) and 698 64×64 images of a statue with varying pitch and yaw angles, as well as lighting (Figure 9, right). Each image can be considered as a point on the “visual manifold” embedded in a high-dimensional space (12K-dimensional in the case of faces and 4096-dimensional in case of statues); the structure and intrinsic dimension of the two manifolds is different (the statue has additional dimensions due to varying lighting). In our discrete setting, the intrinsic structure of the manifolds was represented as unnormalized graph Laplacians (1) with 20-nearest neighbors and self-tuning scale [52]. 25 pairs of roughly corresponding face and statue poses shown in Figure 9 were used as




Figure 8: Examples of correspondence between near-isometric SCAPE human shapes (first row); TOSCA shapes contaminated with geometric and topological noise (second row); non-isometric shapes (third row); and mesh to point cloud (fourth row). Leftmost shape is used as reference. Similar colors encode corresponding points. (Compare to Figures 3, 4, and 6 in [35])

the data (\mathbf{F}, \mathbf{G}) . The goal of the experiment was to establish a meaningful correspondence between the two visual manifolds given this sparse set of corresponding points. We compared the functional correspondences obtained with the methods of Ovsjanikov et al. [31], Pokrass et al. [35], and the proposed method, for which we used the settings $\mu_1 = \mu_2 = 10^{-8}$, $\mu_3 = 10^{-8}$, $\mu_4 = 5 \cdot 10^{-3}$, $k' = 18$, and $k = 10$.

Figure 10 depicts the obtained correspondences, showing that our method achieves a better correspondence in terms of the proximity to groundtruth pitch and yaw angles.

Flickr tagged images We used the multimodal dataset from [14], a challenging subset of 145 annotated Flickr images taken from the NUS-WIDE dataset [11]. The visual data is represented as 64-dimensional color histograms; the text annotations are represented by 1000-dimensional bags of words. The dataset is noisy and heterogeneous: each image has 3–29 tags some of which are irrelevant or wrong.

We represented the image and text modalities structures using graph Laplacians constructed using 20 and 17 nearest neighbors, respectively, and self-tuning scale [52]. 25 corresponding functions were used as the data to construct the functional correspon-




Figure 9: Examples of corresponding human face (left) and statue (right) images. Pairs of corresponding images (points on the respective visual manifolds) were used to find the functional maps between the two visual manifolds.



Figure 10: Correspondences obtained with different methods. Numbers are the coordinates of the face image in the 2-dimensional normalized pose space. From top to bottom: original man face images (first row), ground-truth correspondence statue images (second row), Ovsjannikov et al. [31] (third row), Pokrass et al. [35] (fourth row), and the proposed method (last row). Correspondence closest to the ground truth angles are denoted in bold blue.




Figure 11: Example of image annotations obtained with our method, represented as word clouds (font size is proportional to the tag significance).

dence using our method with settings $\mu_1 = \mu_2 = 10^{-1}$, $\mu_3 = 1$, $\mu_4 = 2 \cdot 10^{-1}$, $k' = 100$, and $k = 55$. Then, for each image we created a word cloud as follows: the absolute values of the functional map of the image on the tags manifold was used as the weight; seven most important tags were selected and shown in font size proportional to the tag weight (Figure 11). Note that since some images have less than 3 groundtruth tags, hence the rest of the tags may be irrelevant; however, their importance is tiny (e.g. rightmost image in Figure 11).

5 Conclusions

We presented a novel method for finding functional correspondence between manifolds based on the geometric matrix completion framework. We model the correspondence between the spaces of functions on two manifolds as a matrix, whose rows and columns have the geometric structure of the underlying manifolds, captured through their Laplacian operators. We discuss several flavors of our problem, and present an efficient subspace version thereof.

Experimental results on standard shape correspondence benchmarks show that our method beats some of the recent state-of-the-art methods, and compares particularly favorably to previous functional correspondence methods in situations when the given correspondence data is very scarce. This allows tackling challenging cases such as matching non-isometric shapes, shapes with missing parts, different representation (meshes vs point clouds) and shapes corrupted by geometric and topological noise. Furthermore, we applied our method to general high-dimensional multimodal data and the obtained correspondences are meaningful and robust. The proposed approach may cope with any other multimodal data – the only change would be the construction of the Laplacian that captures its geometric structure.

In terms of optimization, while `manopt` is well-suited for smooth optimization, it would be advantageous to devise a version of manifold optimization for non-smooth functions such as the L_1 -norm. An alternative path is to use Augmented Lagrangian methods.

On a more philosophical note, we believe that the importance of our work is in bringing a well-known signal processing technique into the domain of geometric processing, continuing the trend of [35] who introduced sparse coding to the problem of

shape correspondence. Conversely, our work follows the very recent trend of bringing geometric structure into signal processing problems such as compressed sensing and matrix completion. We believe that a cross-fertilization of these two fields could be very fruitful.

References

- [1] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. Scape: shape completion and animation of people. *TOG*, 24(3):408–416, 2005.
- [2] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. 4DMOD*, 2011.
- [3] P. J. Besl and N. D. McKay. A method for registration of 3D shapes. *PAMI*, 14(2):239–256, 1992.
- [4] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *JMLR*, 15:1455–1459, 2014.
- [5] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- [6] A. Bronstein, M. Bronstein, and R. Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008.
- [7] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS*, 103(5):1168–1172, 2006.
- [8] A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro. A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *IJCV*, 89(2-3):266–286, 2010.
- [9] M. M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proc. CVPR*, 2010.
- [10] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. In *Proc. Conf. Robotics and Automation*, 1991.
- [11] T. S. Chua, J. Tang, R. Hong, H. Li, Z. Luo, and Y. T. Zheng. Nus-wide: A real-world web image database from national university of singapore. In *Proc. CIVR*, 2009.
- [12] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *PNAS*, 102(21):7426–7431, 2005.
- [13] A. Elad and R. Kimmel. Bending invariant representations for surfaces. In *Proc. CVPR*, 2001.
- [14] D. Eynard, K. Glashoff, M. M. Bronstein, and A. M. Bronstein. Multimodal diffusion geometry by joint diagonalization of laplacians. *arXiv:1209.2295*, 2012.
- [15] J. Ham, D. Lee, and L. Saul. Semisupervised alignment of manifolds. In *Proc. Workshop on Artificial Intelligence and Statistics*, 2005.
- [16] O. Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 20:1–23, 2010.
- [17] V. Kalofolias, X. Bresson, M. Bronstein, and P. Vandergheynst. Matrix completion on graphs. *arXiv:1408.1717*, 2014.
- [18] V. G. Kim, Y. Lipman, X. Chen, and T. A. Funkhouser. Möbius transformations for global intrinsic symmetry analysis. *Computer Graphics Forum*, 29(5):1689–1700, 2010.

- [19] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. *TOG*, 30(4):79, 2011.
- [20] A. Kovnatsky, M. M. Bronstein, A. M. Bronstein, K. Glashoff, and R. Kimmel. Coupled quasi-harmonic bases. *Computer Graphics Forum*, 32:439–448, 2013.
- [21] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. In *Proc. ICCV*, 2005.
- [22] B. Lévy. Laplace-Beltrami eigenfunctions towards an algorithm that understands geometry. In *Proc. SMI*, 2006.
- [23] Y. Lipman and I. Daubechies. Conformal Wasserstein distances: Comparing surfaces in polynomial time. *Advances in Mathematics*, 227(3):1047 – 1077, 2011.
- [24] R. Litman, A. M. Bronstein, and M. M. Bronstein. Diffusion-geometric maximally stable component detection in deformable shapes. *Computers & Graphics*, 35(3):549–560, 2011.
- [25] D. Mateus, R. P. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proc. CVPR*, 2008.
- [26] F. Mémoi. Gromov-Wasserstein Distances and the Metric Approach to Object Matching. *Foundations of Computational Mathematics*, pages 1–71, 2011.
- [27] F. Mémoi and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5(3):313–347, 2005.
- [28] M. Meyer, M. Desbrun, P. Schröder, and A. H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization and Mathematics III*, pages 35–57, 2003.
- [29] N. J. Mitra, N. Gelfand, H. Pottmann, and L. J. Guibas. Registration of point cloud data from a geometric optimization perspective. In *Proc. Eurographics Symposium on Geometry Processing*, pages 23–32, 2004.
- [30] T. Neumann et al. Compressed manifold modes for mesh processing. *Computer Graphics Forum*, 33(5):35–44, 2014.
- [31] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: A flexible representation of maps between shapes. *TOG*, 31(4), 2012.
- [32] M. Ovsjanikov, Q. Mérigot, F. Mémoi, and L. Guibas. One point isometric matching with the heat kernel. *Computer Graphics Forum*, 29(5):1555–1564, 2010.
- [33] V. Ozoliņš, R. Lai, R. Caflisch, and S. Osher. Compressed modes for variational problems in mathematics and physics. *PNAS*, 110(46):18368–18373, 2013.
- [34] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2:15–36, 1993.
- [35] J. Pokrass et al. Sparse modeling of intrinsic correspondences. *Computer Graphics Forum*, 32:459–468, 2013.
- [36] E. Rodolà, A. M. Bronstein, A. Albarelli, F. Bergamasco, and A. Torsello. A game-theoretic approach to deformable shape matching. In *Proc. CVPR*, 2012.
- [37] E. Rodolà, S. R. Buló, T. Windheuser, M. Vestner, and D. Cremers. Dense non-rigid shape correspondence using random forests. In *Proc. CVPR*, 2014.
- [38] Y. Sahillioğlu and Y. Yemez. Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Computer Graphics Forum*, 30(5):1461–1470, 2011.
- [39] J. Shotton, T. Sharp, A. Kipman, A. Fitzgibbon, M. Finocchio, A. Blake, M. Cook, and R. Moore. Real-time human pose recognition in parts from single depth images. *Communications of the ACM*, 56(1):116–124, 2013.
- [40] A. Shtern and R. Kimmel. Matching lbo eigenspace of non-rigid shapes via high order statistics. *arXiv:1310.4459*, 2013.
- [41] P. Skraba, M. Ovsjanikov, F. Chazal, and L. Guibas. Persistence-based segmentation of deformable shapes. In *Proc. NORDIA*, 2010.

- [42] J. Solomon, A. Nguyen, A. Butscher, M. Ben-Chen, and L. Guibas. Soft maps between surfaces. In *Computer Graphics Forum*, volume 31, pages 1617–1626, 2012.
- [43] N. Srebro, J. Rennie, and T. S. Jaakkola. Maximum-margin matrix factorization. In *Proc. NIPS*, 2004.
- [44] A. Tevs, A. Berner, M. Wand, I. Ihrke, and H.-P. Seidel. Intrinsic shape matching by planned landmark sampling. *Computer Graphics Forum*, 30(2):543–552, 2011.
- [45] L. Torresani, V. Kolmogorov, and C. Rother. Feature correspondence via graph matching: Models and global optimization. In *Proc. ECCV*, 2008.
- [46] S. Umeyama. An eigendecomposition approach to weighted graph matching problems. *PAMI*, 10(5):695–703, 1988.
- [47] B. Vandereycken. Low-rank matrix completion by Riemannian optimization. *SIAM J. Optimization*, 23(2):1214–1236, 2013.
- [48] U. Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.
- [49] C. Wang and S. Mahadevan. Manifold alignment without correspondence. In *IJCAI*, 2009.
- [50] F. Wang, Q. Huang, M. Ovsjanikov, and L. Guibas. Unsupervised multi-class joint image segmentation. In *Proc. CVPR*, 2014.
- [51] A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud. Surface feature detection and description with applications to mesh matching. In *Proc. CVPR*, 2009.
- [52] L. Zelnik-Manor and P. Perona. Self-tuning spectral clustering. In *Proc. NIPS*, 2004.
- [53] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios. Dense non-rigid surface registration using high-order graph matching. In *Proc. CVPR*, 2010.